jamf

JAMF'S AI GOVERNANCE SURVEY:

What 687 IT and Security Leaders Revealed About Governing AI

Summary of findings

687 IT and security leaders across Apple enterprises provided insight on the extent, goals and security of their AI deployment. This is what we found.

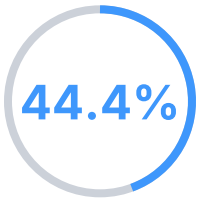

Automating

Deploying

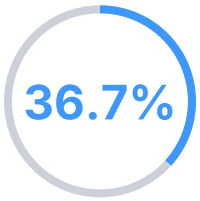

Governing

Three AI priorities converge

Respondents named automating IT operations (44.4%), deploying AI productivity tools (41.0%) and establishing AI governance (36.7%) **as their top AI priorities**.

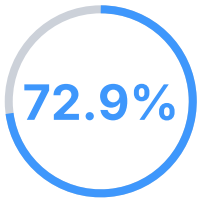

of organizations have deployed AI

Nearly three-quarters of organizations have deployed AI in some form. We're past the point of deciding adoption. Governance is a must.

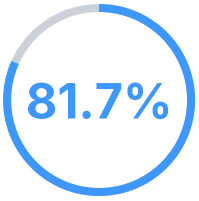

of organizations are exposed to AI risk

22.0% have already had a cost or security incident. Another 59.7% see one as a near-term risk. AI risk is either already lived or actively expected.

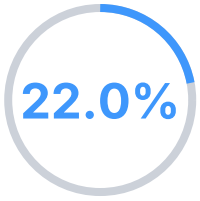

of organizations have had a cost or security incident

Over 1 in 5 organizations have already had an incident related to cost, security or both. The impact lands on the budget and the security team at the same time.

increase in incident rate

Among organizations with deeply integrated AI, 27.1% have had an AI-related incident, compared to 19.4% of those still exploring. Exposure scales with adoption, not against it.

The greater your AI use, the greater your risk.

Most organizations are pushing AI adoption, but risk scales with it. Shadow AI, quietly embedded software features and on-device and agentic tools create blind spots that are hard to govern and harder to audit. As adoption deepens, so does exposure. The question is no longer whether an incident will happen, but when.

CHART 1

Where Apple enterprises are with AI adoption

Key takeaway: Nearly three-quarters of Apple enterprises have deployed AI in some form, from team-level pilots to deep integration into daily workflows. The conversation has moved past whether to adopt.

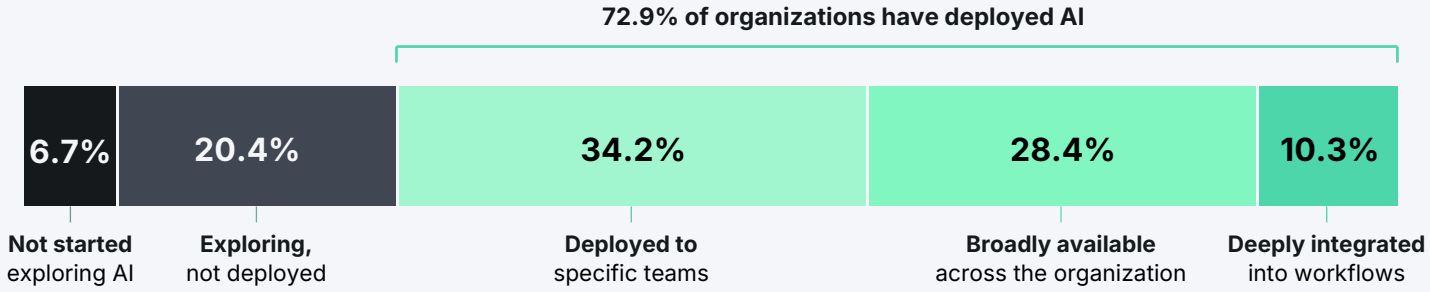

Footnote: n = 687 IT and security leaders. Q2 2026.

Even though nearly 3 in 4 organizations have deployed AI, deployment at scale doesn't reduce risk. It compounds it. 22.0% have already experienced an incident: 7.3% from unexpected cloud AI or LLM costs, 5.1% from a security or data exposure event, and 9.5% from both. Of those that haven't, 59.7% still expect one.

CHART 2

AI-related incidents and concerns over the past 12 months

Key takeaway: 22.0% of organizations have already had an AI-related incident. Another 59.7% expect one. Only 18.3% say AI is not currently a concern.

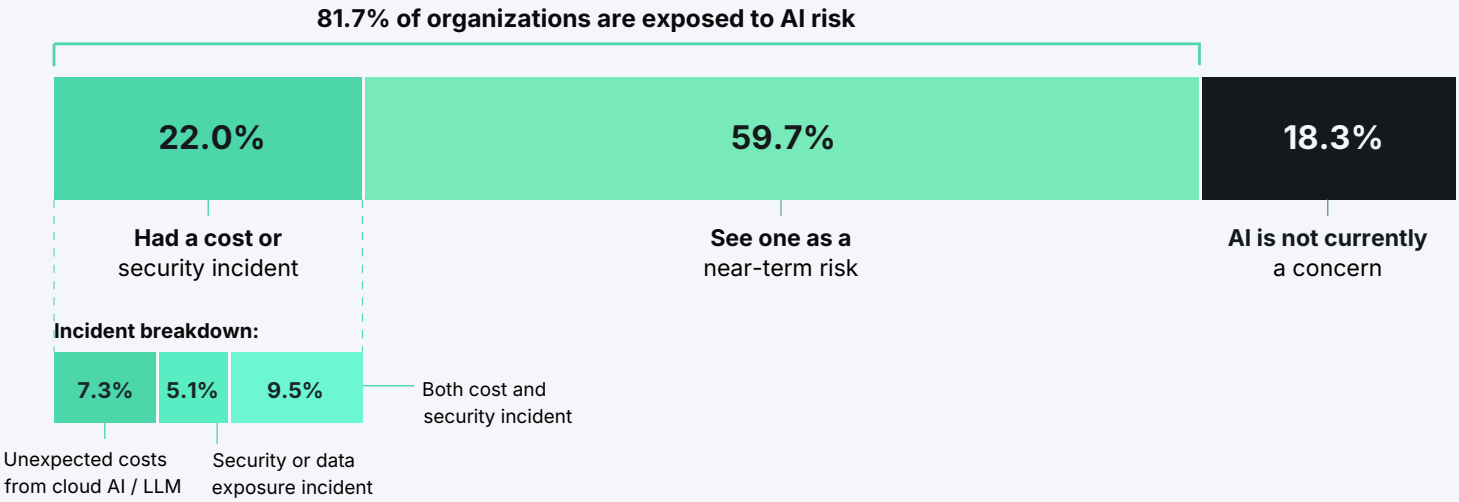

Footnote: *n = 681. Pattern holds in both survey samples independently.*

However, the data holds a counterintuitive finding: the organizations furthest along with AI report incidents at the highest rate.

CHART 3

Incident rate by AI adoption depth

Key takeaway: Apple enterprises with deeper AI deployment report incidents at higher rates. Among the most mature adopters, 27.1% have already had an incident, compared to 19.4% of those still exploring.

Footnote: n = 683 IT and security leaders. Q2 2026.

Once teams begin exploring AI in their organization, the chance they've had an incident rises. The share of organizations reporting an incident is 40% higher among those with deeply integrated AI (27.1%) than among those still exploring (19.4%).

AI challenges follow common themes

Respondent's free-response answers about incident prevention clustered around four themes.

Shadow AI

Boosts to productivity, a boom in AI tools and company-wide pushes for AI integration mean that employees are turning to AI on a regular basis. This often happens without IT approval; employees create personal accounts and may input sensitive data. As a result, IT is left in the dark about what AI systems are used, making AI platforms hard to control or block. A lack of visibility makes security and governance difficult, if not impossible.

Vendor sprawl

Beyond the surge of new AI-based software, many apps are pushing AI into their pre-existing products. Vetting and deploying each possible AI tool is time consuming and difficult for IT teams, especially at the speed AI is moving. Respondents note difficulty deciding what AI platforms are best for their employees and pushing employees to use sanctioned AI tools. These increasing points of entry make AI difficult to secure.

Agentic and developer AI

Challenges with agentic and developer AI show up in a few key areas: secure deployment/visibility, AI features and user education. Respondents cite issues managing agentic AI deployment in a way that enables users without putting data at risk. Problems with visibility into command-line tools, third-party packages, IDE extensions, embedded LLMs and more are also common. With appropriate permissions, agentic AI opens serious risks to code bases if insecure or problematic code is added or necessary code is removed. Development issues extend to users that aren't developers themselves, as they create their own apps without proper vetting and quality checks.

Cost surprises

Balancing costs, company initiatives and security strains IT teams. Usage-based pricing on cloud AI and LLM APIs makes spend hard to forecast, and as departments rapidly adopt new tools, overlapping paid licenses accumulate. Without visibility into what's actually used, IT teams have no clear way to decide which tools to consolidate.

Governance and productivity go hand in hand.

Of AI adopters, the more deeply AI is embedded, the greater the incident rate.

Respondents' common challenges span visibility, deployment, vendor sprawl and cost.

Together, these findings make clear a single reality: **AI is being implemented faster than it can be governed.**

This manifests in shadow AI, exposed entry points into company data or systems, redundant (and costly) platform adoption, and risks that are hard to measure because they're hidden from IT.

In response, IT teams need to adjust their priorities now that AI reshapes how everyone works.

CHART 4

Top AI priorities for the next 12 months

Key takeaway: Automating IT operations, deploying productivity tools, and establishing governance are pursued at similar levels. Governance is not trailing enablement, but moving along with it.

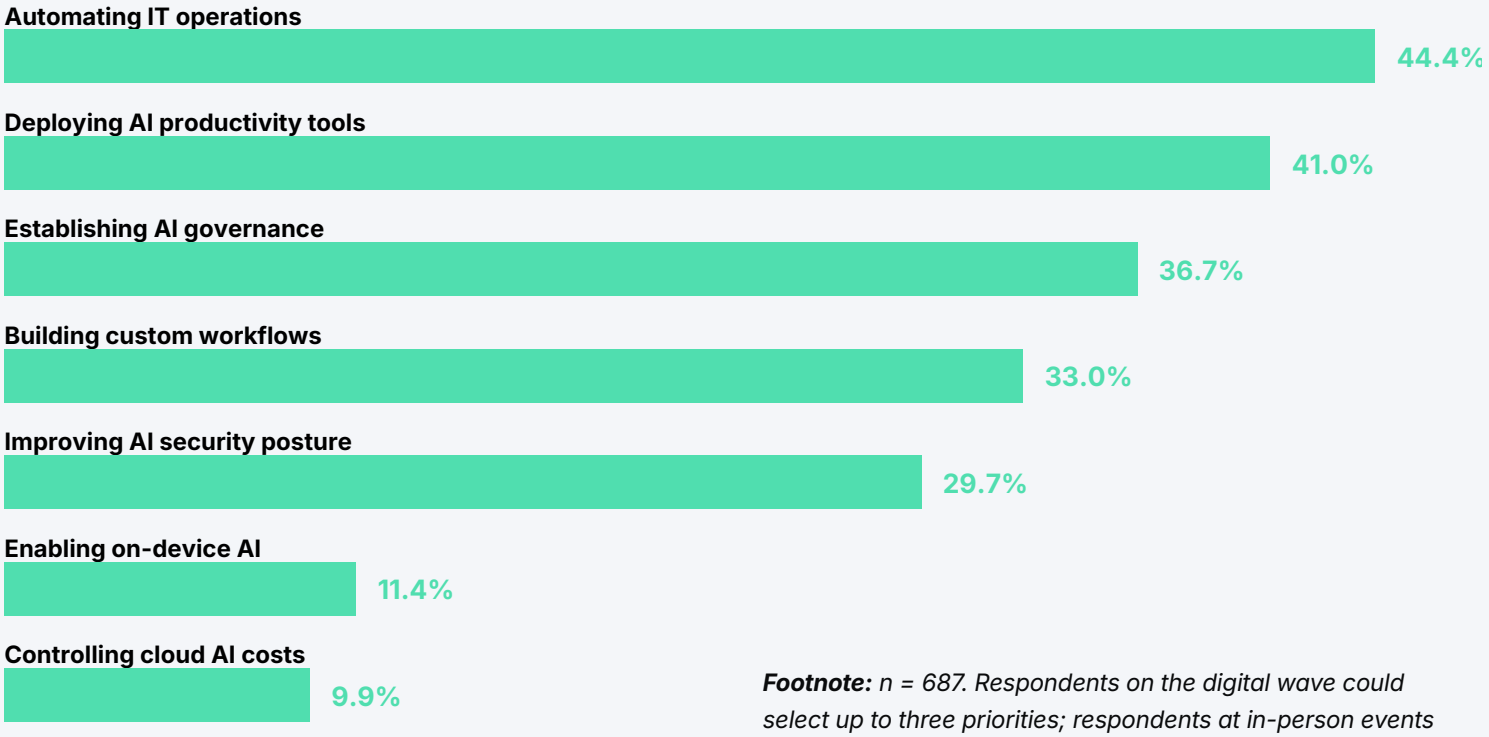

Footnote: *n = 687. Respondents on the digital wave could select up to three priorities; respondents at in-person events selected one. See methodology for details.*

Governance and enablement may seem like opposing processes. The more AI tools there are, the harder they are to govern. IT teams have always had to balance competing priorities, and AI is no exception. The speed AI deployment moves at; its features and risks cross into new territory.

That's why teams are pursuing these priorities at the same time. If you move too fast, you increase your chance of an incident. Move too slow, and employees find workarounds that hinder your security posture.

💬 In your own words: AI challenges faced by your peers

In 178 open-ended responses, the survey gained detailed insights that contextualize the findings. Eight voices from the open-ended responses we received*:

Users want access now, and security teams pushing back are feeling the pressure. The core problem remains unsolved: **full control tends to hurt productivity, but relaxing controls opens real compliance risk.**

Every vendor is embedding AI whether you asked for it or not. Blanket disabling buys time but isn't sustainable. The bigger concern is how data is processed in the cloud and whether we can control where our data goes.

Blocking known AI sites is the easy part. CLI tools, IDE extensions, browser extensions, and packages pulled from GitHub **are largely invisible**, and when one vector gets closed, users find another.

In regulated industries and certain jurisdictions, specific compliance frameworks have to be in place before anything gets turned on, and right now the **tools and frameworks aren't meeting requirements.**

Shadow AI and black-box script execution top the risk list. Close behind is non-technical users vibe-coding their own apps, building things they don't fully understand, with data exposure they're not aware of.

There's a gap between the expectation that everyone uses AI and the willingness to fund the licenses. Teams that moved fast now have multiple overlapping agents with **high costs and no clear framework** for deciding which tools are worth keeping.

Giving AI agents access to dev and prod infrastructure is a hard sell. The fear is an agent that does exactly what you didn't want, then tells you the data is gone. **Deploying agentic capabilities in a controlled, managed way is still an unsolved problem.**

When hallucinated outputs get treated as fact, and AI keeps embedding itself deeper into daily routines, the **risks compound faster than the literacy does.**

* The themes above were composed by Jamf based on patterns across 178 open-ended responses. Each captures a recurring theme rather than the verbatim words of a single respondent.

Take action: four governance principles

1.

Gain visibility.

As many respondents explained, gaining visibility is paramount. You can't govern what you can't see. But, of course, therein lies the difficulty. Frequent auditing of installed apps and traffic monitoring will help identify interactions with AI platforms. As users leverage local AI platforms and already approved non-AI applications add AI into their feature sets, deeper investigation into AI runtime detection is required.

2.

Govern the tool, not the user.

For many IT teams, their organization's AI policies came quickly and without IT consideration. And these policies encourage more AI use, as soon as possible. Even if user guidance is provided, guidance isn't enforcement. This is where shadow AI emerges.

Instead, governance must be decided based on the organization's risk tolerance and security guidelines, which should be reflected in AI tool's data sharing settings: what data it accesses, how it handles it and what it can change. With shadow AI, the user isn't always visible, but the traffic, data and API calls are. You can only govern what you can see.

3.

Build governance into deployment.

Organizations that deployed AI too quickly exposed themselves to possible incidents. Order matters — governance must accompany app deployment, not react to it. Yes, this is easier said than done, and you may already be working against a backlog. But by identifying what tools are being used, offering them to users and establishing access policies will make it easier to securely scale your AI tools.

4.

Use tools built for Apple, not bolted onto it.

Network-based tools show you the traffic: which cloud AI services users hit, when, and how often. That's a real signal, but it stops at the network edge. Even when the AI itself runs in the cloud, the access happens on the device: which tools are installed, what processes they spawn, and what files they access. None of that shows up in a DNS log. Apple-native tooling closes the gap: see the tools, the processes, the file access, and enforce which ones are allowed.

Govern what you enable. Enable what you govern.

AI moves faster than most governance frameworks were built to handle. But the challenges you've already faced don't have to define how you deploy from here. There's no need to choose between giving users the AI tools they need and securing how those tools are used.

The teams pulling ahead aren't moving fastest or locking down hardest. They're treating governance and enablement as the same project, building visibility and access controls into AI deployment from the start. For Apple enterprises, that depends on tooling that understands the runtime you're managing: cloud traffic, on-device models and agentic processes each leave different signals, and monitoring not built for Apple misses most of them. The governance you build is only as strong as what your tools can see.

You can't slow down AI adoption. But you can govern it — and that's where the work starts.

Methodology

The data was collected in two waves. Wave one was fielded to the Jamf customer community in March and April 2026 (338 respondents). Wave two was an in-person survey at Jamf Nation Live events across six North American cities (349 respondents). Combined respondent base: 687 IT and security leaders. All respondents work in organizations that manage and secure Apple devices at scale, as Jamf customers.

On the priorities question, respondents were asked to select their top AI priorities for the next 12 months. The March and April surveys allowed up to three selections; the Jamf Nation Live survey allowed one. Percentages on this question represent the share of all 687 respondents who selected each priority among their top selections. Both selection rules are disclosed here for transparency.

Statistical testing confirms that the two waves produced distinct populations, with Jamf Nation Live respondents on average earlier in their AI maturity. The directional findings hold in both samples independently. All respondent data was collected and analyzed on an anonymized basis; no individual responses are attributed to specific respondents or organizations. Respondents were not compensated for participation.